

The Oracle logo is displayed in a red, sans-serif font. It is positioned in the upper left area of the slide, partially overlaid by decorative abstract shapes: a red and black patterned shape on the far left and a blue shape with green wavy lines above it.

ORACLE

Photo by dreokt on Unsplash

Migration to Oracle Autonomous Database

Part 4: The DBA's Life

Oracle

DBAs

run the world





MIKE DIETRICH

Vice President

Database Upgrade, Migrations & Patching

 mikedietrich

 @mikedietrichde.com

 <https://mikedietrichde.com>





DANIEL OVERBY HANSEN

Distinguished Product Manager
Database Upgrade, Migrations & Patching

 dohdatabase

 @dohdatabase.com

 <https://dohdatabase.com>



RODRIGO JORGE

Distinguished Product Manager
Database Upgrade, Migrations & Patching

 [rodrigoaraujorge](#)

 [@dbarj.com.br](#)

 <https://dbarj.com.br>



ALEX ZABALLA

Distinguished Product Manager
Database Upgrade, Migrations & Patching



alexzaballa



@alexzaballa.bsky.social



<https://alexzaballa.com>




Find Slides and Much More on Our Blogs



MikeDietrichDE.com

Mike.Dietrich@oracle.com



dohdatabase.com

Daniel.Overby.Hansen@oracle.com



DBArj.com.br

Rodrigo.R.Jorge@oracle.com




AlexZaballa.com

Alex.Zaballa@oracle.com



Download the Slides

<https://MikeDietrichDE.com/slides>





<https://MikeDietrichDE.com/videos>

More than 40 hours of technical content,
on-demand, anytime, anywhere





Patching



Oracle patches your Autonomous Database

- Patching happens in pre-defined maintenance windows

Autonomous Database Serverless



Oracle has a two-hour
maintenance window every weekend

- Open an SR to change the period



```
select    expected_start_date, maintenance_product,  
          maintenance_status, patch_id  
from      db_notifications;
```

EXPECTED_START_DATE	MAINTENANCE_PRODUCT	MAINTENANCE_STATUS	PATCH_ID
-----	-----	-----	-----
2025-05-04T02:00:00Z	INFRASTRUCTURE	COMPLETED	ADBS-25.4.4.2
2025-05-11T02:00:00Z	DICTIONARY	COMPLETED	ADBS-25.4.5.2
2025-05-25T02:00:00Z	DATABASE	COMPLETED	ADBS-25.5.2.2
2025-06-01T02:00:00Z	DICTIONARY	COMPLETED	ADBS-25.5.4.2
2025-06-08T02:00:00Z	DATABASE	COMPLETED	ADBS-25.5.5.2
2025-06-15T02:00:00Z	DICTIONARY	SCHEDULED	null





Maintenance Products

Infrastructure

Grid Infrastructure and operating system

Entire Exadata stack including firmware

Less frequent cadence

Database

Rolling maintenance

Sessions are killed if exceeding drain timeout

Dictionary



Maintenance Products

Infrastructure

Database

Dictionary

Database Oracle home

2-week cadence

Rolling maintenance

Sessions are killed if exceeding drain timeout



Maintenance Products

Infrastructure

Datapatch on CDB\$ROOT

Database

Capture/Replay on PDBs

```
alter pluggable database application  
APP$CDB$DATAPATCH sync
```

Dictionary

2-week cadence

Sessions are killed if they use objects up for replacement





Autonomous Database Serverless does not use the regular Release Update



```
select component_name, bug_num, bug_title,  
from   dba_cloud_patch_info  
where  patch_version='ADBS-25.5.5.2';
```

COMPONENT_NAME	BUG_NUM	BUG_TITLE
-----	-----	-----
RDBMS	36716948	AIM:ORA-600 [KRCPPSP_NOT_PRIMARY] - KRCPPSP
METADATA	37807168	OCI:: PHX::PHX:E62POD BUG 37805388 - PHX E62POD UNABLE ...
METADATA	38012226	SET SESSION MULTIPLIER TO 2X IN FA PDBS AND ALSO SET ...
METADATA	37939748	ADD A DRY-RUN FEATURE IN ADB MIGRATION
RDBMS	37554264	AZURE:: IAD::IAD:E79POD IMPORT FAILURE WITH WORKERS ...
RDBMS	37379354	INSERT ALL INTO THE TABLE WITH A TRIGGER FAILS WITH ...
.		
.		
.		



Oracle applies critical security patches
as soon as they are available

- Even outside predefined maintenance windows



Oracle validates patches in multiple stages: *Early and Regular*


- Always Free is in stage *Early*

Patch Stage




- Not for customers
- Several internal stages

Patch Stage



- Always Free
- Production-ready patches
- First stage for customers

Patch Stage





Oracle monitors all stages for errors

- Feeds into Automatic Incident Monitoring

BUG 32765738 - AIM:ORA-7445 [KEWSSYSVF_1] - KEWSSYSVF_1

[Edit this bug](#) [Bug Tree](#) [ARU Re](#)

Bug No: 32765738

Filed By:

Updated By:

Sup Rep:

Customer:

Status: 80 - Development to QA/Fix Delivered Internal

SRs in Bug Hierarchy: Open: 12 Others: 274 ([Show Details](#))

Change the patch level

Patch level

☐ Regular

☒ Early

Warning

The time it will take to change the patch level depends on the size of your database. You may see brief connection drops during this time. [Learn more](#)

Patch Level: Idea



EARLY

Dev, Test

REGULAR

Pre-prod, Prod



File an SR if you find a regression
in your *Early* patch level environments

- [Zero-Regression Service Level Objective](#)

Testing Patches




EARLY

Dev, Test, Patch-test

REGULAR

Pre-prod, Prod

Testing Patches



- Refreshable clone PDB
- Capture workload on Prod
- Replay workload on Patch-test
- Automatically repeated

```
--First, define a workload
BEGIN
    DBMS_CLOUD_ADMIN.ENABLE_FEATURE(
        feature_name => 'WORKLOAD_AUTO_REPLAY',
        params       => JSON_OBJECT(
            'target_db_ocid' VALUE '...',
            'capture_duration' VALUE 120,
            'capture_day' VALUE 'MONDAY',
            'capture_time' VALUE '15:00'));

END;
/
```

```
--First, define a workload
BEGIN
    DBMS_CLOUD_ADMIN.ENABLE_FEATURE(
        feature_name => 'WORKLOAD_AUTO_REPLAY',
        params       => JSON_OBJECT(
            'target_db_ocid' VALUE '...',
            'capture_duration' VALUE 120,
            'capture_day' VALUE 'MONDAY',
            'capture_time' VALUE '15:00'));

END;
/

--Next, check the results after it has been replayed
SELECT * FROM DBA_CAPTURE_REPLAY_STATUS;
```

Autonomous Database Dedicated



You specify a 4-hour quarterly maintenance window

- You can skip a quarter, but not two consecutive quarters





When needed, Oracle schedules monthly infrastructure security patching

- Vulnerabilities with CVSS scores greater than or equal to 7



Maintenance Products

Quarterly

Patches Exadata Infrastructure, VM Cluster,
and Container Database

Includes Release Updates
- Rolling or non-rolling

Monthly

Specify a buffer time between standby and primary

Optionally, also timezone file upgrades

One-offs

Edit automatic maintenance

You either choose a schedule for quarterly automatic maintenance of the Autonomous Container Database (ACD), or the update will be automatically scheduled for your ACD. Oracle notifies you, in advance of the date and time, of upcoming scheduled maintenance. You can change the schedule at any point before the beginning of the scheduled maintenance. Scheduled maintenance for subsequent quarters does not affect the schedule for the current quarter. [Learn more](#)

Configure maintenance method

Rolling

The Autonomous Container Database is updated one node at a time with no downtime for the Autonomous Databases.

Non-rolling

The Autonomous Container Database is shut down and updated in parallel across all nodes. This method minimizes maintenance time but requires a full Autonomous Container Database and all associated Autonomous Databases downtime.

Enable time-zone update



Time-zone files can only be updated using the non-rolling configuration method.

Configure Autonomous Container Database maintenance version

Next release update (RU)

Update to next RU during next scheduled maintenance.

Latest release update (RU)

Update to latest RU during next scheduled maintenance.

Configure automatic maintenance schedule

No preference

System assigns a start time and date for Autonomous Container Database maintenance.

Custom schedule

Choose a preferred month, week, weekday, and start time for Autonomous Container Database maintenance.



ExaCS-X8M

Available

Exadata Infrastructure

Actions

Scale Infrastructure

An Exadata Infrastructure (EI) resource represents a physical Exadata Database system and is composed of an elastically-expandable set of Compute and Storage servers. It hosts the underlying components Autonomous VM Clusters, Autonomous Container Databases, and Autonomous Databases.

Exadata Infrastructure Information

Maintenance

Autonomous Exadata VM clusters

DB servers

Maintenance history

Work requests

Tags

Title	State	Target Database Server...	Target storage server...	Start time (UTC)	End time	Type	Scheduling	
Cloud Exadata Infrastructure Maintenance	Succeeded	25.1.4.0.0.250512	25.1.4.0.0.250512	Sun, Jun 1, 2025, 00:00:36 UTC	Sun, Jun 1, 2025, 03:27:10 UTC	Quarterly	Planned	...
Cloud Exadata Infrastructure Security Maintenance	Succeeded	250504	24.1.11.0.0.250413	Tue, May 27, 2025, 21:30:07 UTC	Tue, May 27, 2025, 23:41:52 UTC	Monthly security	Planned	...
Cloud Exadata Infrastructure Security Maintenance	Succeeded	250504	24.1.10.0.0.250513	Mon, Apr 28, 2025, 21:30:07 UTC	Mon, Apr 28, 2025, 23:37:05 UTC	Monthly security	Planned	...
Cloud Exadata Infrastructure Security Maintenance	Succeeded	250504	24.1.9.0.0.250221	Tue, Apr 8, 2025, 20:30:08 UTC	Tue, Apr 8, 2025, 22:37:27 UTC	Monthly security	Planned	...
Cloud Exadata Infrastructure Security Maintenance	Succeeded	250109.1	24.1.8.0.0.250208	Wed, Mar 5, 2025, 20:30:09 UTC	Wed, Mar 5, 2025, 22:37:25 UTC	Monthly security	Planned	...
Cloud Exadata Infrastructure Security Maintenance	Succeeded	250109.1	24.1.7.0.0.241204	Wed, Feb 5, 2025, 01:00:11 UTC	Wed, Feb 5, 2025, 01:35:48 UTC	Monthly security	Planned	...
Cloud Exadata Infrastructure Maintenance	Succeeded	24.1.8.0.0.250130	24.1.8.0.0.250130	Sun, Feb 2, 2025, 01:00:33 UTC	Sun, Feb 2, 2025, 04:05:58 UTC	Quarterly	Planned	...
Cloud Exadata Infrastructure Maintenance	Succeeded	24.1.5.0.0.241115	24.1.6.0.0.241115	Sat, Dec 28, 2024, 00:00:34 UTC	Sat, Dec 28, 2024, 01:55:50 UTC	Quarterly	Planned	...
Cloud Exadata Infrastructure Security Maintenance	Succeeded	241105.1	24.1.6.0.0.241115	Mon, Nov 25, 2024, 00:00:06 UTC	Mon, Nov 25, 2024, 02:13:08 UTC	Monthly security	Planned	...
Cloud Exadata Infrastructure Security Maintenance	Succeeded	241005	24.1.5.0.0.241016	Sun, Nov 3, 2024, 21:30:07 UTC	Mon, Nov 4, 2024, 00:02:41 UTC	Monthly security	Planned	...





Find information about quarterly maintenance on My Oracle Support

- [Exadata Database Machine and Exadata Storage Server Supported Versions \(Doc ID 888828.1\)](#)



Maintenance Products

Quarterly

Monthly

One-offs

At least 7 days notice

Option to re-schedule within a period

Not possible to skip

No database interruption, but you can't perform certain administrative tasks



Maintenance Products

Quarterly

Monthly

One-offs

Request one-offs via a service request

Oracle applies the one-off within 72 hours of it being available

You may choose to apply it immediate

One-offs are merged into the next Release Update



Autonomous Database Dedicated
does not use the regular Release Update

General Recommendations



To avoid interruptions your application must adhere to the MAA guidelines



Drain timeout is 5 minutes

- Fixed in ADB-S
- Adjustable in ADB-D using `DBMS_APP_CONT_ADMIN.SET_DRAINING`





Align your maintenance window
with an off-peak period



Use OCI Events to receive notification
on maintenance start and finish

Further Reading

ADB-S

- [Patch and Maintenance Windows](#)
- [Use Application Continuity on Autonomous Database](#)
- [Application Checklist for Continuous Service with Autonomous Database on Shared Infrastructure](#)
- [Test Workloads Against an Upcoming Patch](#)

ADB-D

- [Patching](#)
- [Continuous Availability](#)
- [Continuous Availability Best Practices for Applications Using Autonomous Database - Dedicated](#)
- [Exadata Cloud: Frequently Asked Questions about SMR - Monthly Security Maintenance \(Doc ID 2951685.1\)](#)
- [Exadata Database Service on Dedicated Infrastructure - Explanation Of Cloud Operations Service \(Doc ID 2875973.1\)](#)



Real Application Testing



If you don't test, you'll fail - most likely

- Testing requires resources in every projects



Real Application Testing eases your tests

- But it may require a bit of extra work



Real Application Testing | Two Components



SQL Performance Analyzer

Test Execute and verify SQL statements
Optimize individual statements

Real Application Testing



Database Capture/Replay

Replay an entire workload
Measure and optimize performance

SQL Performance Analyzer



Collect workload information
into SQL Tuning Sets (STS)

- Sample from cursor cache
- Gather from AWR



Gather at least a full month of workload data

- Assist in testing your database
- Useful in solving post-migration performance problems

SQL Performance Analyzer



Migrate the database

Load workload data
(SQL Tuning Set)




SQL Performance Analyzer



SQL Performance Analyzer tests your workload

Report all regressing and error statements

SQL Performance Analyzer | Concept



Pro tip: For migrations, import STS into target database

SQL Performance Analyzer | Report

Regressed SQL Statements						
	SQL ID	Net Impact on Workload (%)	Buffer Gets		Net Impact on SQL (%)	New Plan
			SQL Trial 1	SQL Trial 2		
↓	3fv28gfu9y0aq	-0.050	26,504	29,573	-11.580	Y
↓	czzzubf8fjz96	-0.030	1,410	1,981	-40.500	Y

From production
workload

From test
execution

SQL Performance Analyzer | Report

Regressed SQL Statements						
	SQL ID	Net Impact on Workload (%)	Buffer Gets		Net Impact on SQL (%)	New Plan
			SQL Trial 1	SQL Trial 2		
↓	3fv28gfu9y0aq	-0.050	26,504	29,573	-11.580	Y
↓	czzzubf8fjz96	-0.030	1,410	1,981	-40.500	Y

SQL Details: czzzubf8fjz96

Parsing Schema APPSExecution Frequency 3

> SQL Text

SELECT /* my_query_21 */ /*+ ORDERED INDEX(t1) USE_HASH(t1) */ 'B' || t2.take_02 take_02, 'B' || t2.take_15 take_15, 'B' || t2.take_08 take_08, 'x' || t3.record_nr price_eur_id,...

Single Execution Statistics

	Execution Statistic Name	Net Impact on Workload (%)	Execution Statistic Collected		Net Impact on SQL (%)
			SQL Trial 1	SQL Trial 2	
↓	Elapsed Time (sec)	-0.240	0.112	0.364	-46.170
↑	Parse Time (sec)	0.220	0.001	0.001	14.490
↓	CPU Time (sec)	-0.030	0.108	0.114	-5.040
↔	User I/O Time (sec)	0.000	0.000	0.000	0.000
↓	Buffer Gets	-0.030	1,410	1,981	-40.500





SQL Performance Analyzer | Report

Regressed SQL Statements						
	SQL ID	Net Impact on Workload (%)	Buffer Gets		Net Impact on SQL (%)	New Plan
			SQL Trial 1	SQL Trial 2		
↓	3fv28gfu9y0aq	-0.050	26,504	29,573	-11.580	Y
↓	czzzubf8fjz96	-0.030	1,410	1,981	-40.500	Y

Plan Comparison						
SQL_TRIAL_1353942463446						
Plan Hash Value: 1165613724						
Expand All Collapse All						
Operation	Line ID	Object	Rows	Cost	Predicate	
▽ SELECT STATEMENT	0		1	9,830		
▽ HASH GROUP BY	1		1	9,830		
▽ MERGE JOIN	2		1	9,829		
▽ SORT JOIN	3		8	9,795		
▽ HASH JOIN	4		8	9,794	T1".PERIOD_CODE"=T4".FLYER...	
INDEX RANGE SCAN	5	APPS.IDX\$\$_080F0004	1	2	T4".EXPORT_LIC_NR"=14659	
▽ HASH JOIN	6		14,210	9,792	T1".SKU_NR"=T2".SKU_NR" AN...	



SPA | Continuous Improvement



--If your queries have a lot of binds, you may tweak the below
--underscore. It defines how much bind data will be kept in memory
--for each query. Default at 400 is often too low.
--Find procedures and more on:
--<https://blogs.oracle.com/coretec/post/spa-in-autonomous-database>

```
alter system set "_cursor_bind_capture_area_size"=3999;
```

SQL Performance Analyzer



Tune SQLs having regressed plans

Create SQL Plan Baselines or rely on
ADB mechanisms

Database Capture/Replay




Capture Workload
on source environment

Database Capture/Replay



Move capture files to ADB

Database Capture/Replay




Replay workload in ADB

Database Capture/Replay



Monitor results and
optimize performance

Database Capture/Replay



Use Database Capture/Replay for upgrade tests in ADB with Refreshable Clones

[Documentation](#)

Interview



Useful Links

Documentation

- <https://docs.oracle.com/en/database/oracle/oracle-database/23/ratug/introduction-to-oracle-database-testing.html>
- <https://docs.oracle.com/en/cloud/paas/autonomous-database/serverless/adbsb/autonomous-real-application-testing.html#GUID-EB8F065E-5FBB-480D-BAF6-5A0446740073>

Blog Posts

- <https://blogs.oracle.com/coretec/post/spa-und-autonomousdatabase>
- <https://blogs.oracle.com/coretec/post/adb-database-replay>

GIT

- <https://github.com/oracle-devrel/technology-engineering/tree/main/data-platform/core-converged-db/real-application-testing/db-replay>
- <https://github.com/oracle-devrel/technology-engineering/tree/main/data-platform/core-converged-db/real-application-testing/sql-performance-analyzer>



Data Pump

Data Pump Trace

Collect:

- Data Pump log file
- AWR report
- Data Pump trace files
 - Stored in the database trace directory
 - Control process file name: ***dm***
 - Worker process file names: ***dw***



How to trace Data Pump jobs

- MOS Doc ID [286496.1](#)



```
-- Change AWR snap interval to 15 minutes and create snapshot
exec dbms_workload_repository.modify_snapshot_settings(null, 15);
exec dbms_workload_repository.create_snapshot;

-- Run Data Pump job with trace (Doc ID 286496.1)
expdp ... metrics=yes logtime=all trace=1FF0300
impdp ... metrics=yes logtime=all trace=1FF0300

-- Create AWR snapshot and produce AWR report
exec dbms_workload_repository.modify_snapshot_settings(null, <original-value>);
exec dbms_workload_repository.create_snapshot;
```

```
-- Change AWR snap interval to 15 minutes and create snapshot
exec dbms_workload_repository.modify_snapshot_settings(null, 15);
exec dbms_workload_repository.create_snapshot;

-- Run Data Pump job with trace (Doc ID 286496.1)
expdp ... metrics=yes logtime=all trace=1FF0300
impdp ... metrics=yes logtime=all trace=1FF0300

-- Create AWR snapshot and produce AWR report
exec dbms_workload_repository.modify_snapshot_settings(null, <original-value>);
exec dbms_workload_repository.create_snapshot;
```

```
-- Change AWR snap interval to 15 minutes and create snapshot
exec dbms_workload_repository.modify_snapshot_settings(null, 15);
exec dbms_workload_repository.create_snapshot;

-- Run Data Pump job with trace (Doc ID 286496.1)
expdp ... metrics=yes logtime=all trace=1FF0300
impdp ... metrics=yes logtime=all trace=1FF0300

-- Create AWR snapshot and produce AWR report
exec dbms_workload_repository.modify_snapshot_settings(null, <original-value>);
exec dbms_workload_repository.create_snapshot;
```

```
SELECT
  output
FROM
  TABLE(
    DBMS_WORKLOAD_REPOSITORY.AWR_REPORT_HTML(
      l_dbid => YOUR_DBID,
      l_inst_num => YOUR_INSTANCE_NUMBER,
      l_bid => BEGIN_SNAP_ID,
      l_eid => END_SNAP_ID
    )
  );
```

```
SQL> SELECT dbid FROM v$database;
SQL> SELECT instance_number FROM v$instance;
SQL> SELECT snap_id, begin_interval_time, end_interval_time
FROM dba_hist_snapshot ORDER BY snap_id DESC FETCH FIRST 10 ROWS ONLY;
```

```
SQL> SELECT *  
      FROM V$DIAG_TRACE_FILE  
      ORDER BY CHANGE_TIME DESC;
```

```
SQL> SELECT PAYLOAD  
      FROM V$DIAG_TRACE_FILE_CONTENTS  
      WHERE TRACE_FILENAME = 'ef01pod1_ora_267060.trc'  
      ORDER BY LINE_NUMBER ASC;
```



Trace file may be truncated

MAX_DIAG_SIZE defaults to **100M** in ADB





Log file content is spooled to the standard output device of any attached clients



Log file can be skipped

NOLOGFILE=YES command-line option



Still need the log file?

- Print on screen

`UTL_FILE.GET_LINE + DBMS_OUTPUT.PUT_LINE`

- Move and download from Object Storage

`DBMS_CLOUD.PUT_OBJECT`

- Save on your local File System

`SQLcl + JS (java.nio.file.Files.copy)`

Tracing requires permissions



User must be granted any of the following roles:

- EXP_FULL_DATABASE
- IMP_FULL_DATABASE
- DATAPUMP_EXP_FULL_DATABASE
- DATAPUMP_IMP_FULL_DATABASE
- DATAPUMP_CLOUD_EXP
- DATAPUMP_CLOUD_IMP



Open an SR to request a temporary increase

Request "`_pdb_max_diag_size`" to be raised to **10G**

Need more?

Also request:

- `'sql_trace {process: pname = dw | process: pname = dm} level=8'`
- `'39089 trace name context forever, level 0x01ff0f00'`



Tips and tricks



Get easy access to DBA tasks in *Database Actions*





SQL



Data Modeler



REST



Liquibase



JSON



Charts



Scheduling



Machine Learning



APEX



Graph Studio


SQL

The SQL worksheet is where most of your work will be performed. Running SQL queries and scripts, creating and browsing objects, loading data, exporting data to CSV or JSON, and so much more. Your day will often start and end in the SQL interface.



Documentation

[Walk through](#)[Open](#)



SQL



Data Modeler



REST



Liquibase



JSON



Charts



Scheduling



Machine Learning



APEX



Graph Studio

Scheduling

An interface for DBMS_SCHEDULER that enables you to monitor jobs, view execution history, forecast upcoming jobs, and visualize scheduler chains.



Documentation

Open

 Database Users APEX Workspaces Data Pump Download Client Credentials (Wallet) Set Resource Management Rules

Database Users

REST enable schemas, change passwords, assign storage quota, manage schema requests, and identify accounts with passwords that will expire soon.



The screenshot displays the 'Database Users' management interface. On the left, a sidebar lists navigation options: Database Users, APEX Workspaces, Data Pump, Download Client Credentials (Wallet), and Set Resource Management Rules. The main content area is titled 'Database Users' and includes a description: 'REST enable schemas, change passwords, assign storage quota, manage schema requests, and identify accounts with passwords that will expire soon.' Below this, there is a 'Current user' section showing the 'ADMIN' user with a green profile icon and a list of roles. To the right, a 'Create User' dialog is open, showing fields for 'User Name', 'Password', and 'Confirm Password', along with checkboxes for 'Create Schema' and 'Grant Privileges'. At the bottom of the dialog are 'Create User' and 'Cancel' buttons. Below the dialog, a table lists existing database users with columns for 'User Name', 'Role', 'Status', 'Quota', and 'Default Schema'.

User Name	Role	Status	Quota	Default Schema
ADMIN	ADMIN	ACTIVE	UNLIMITED	ADMIN
SCOTT	SCOTT	ACTIVE	100MB	SCOTT
HR	HR	ACTIVE	100MB	HR
SALES	SALES	ACTIVE	100MB	SALES
MARKETING	MARKETING	ACTIVE	100MB	MARKETING
FINANCE	FINANCE	ACTIVE	100MB	FINANCE
OPERATIONS	OPERATIONS	ACTIVE	100MB	OPERATIONS
SYSTEM	SYSTEM	ACTIVE	100MB	SYSTEM
DBSNMP	DBSNMP	ACTIVE	100MB	DBSNMP
SI	SI	ACTIVE	100MB	SI
SI_REPO	SI_REPO	ACTIVE	100MB	SI_REPO
SI_REPO_ADMIN	SI_REPO_ADMIN	ACTIVE	100MB	SI_REPO_ADMIN

[Documentation](#)

Open

 Database Users APEX Workspaces Data Pump Download Client Credentials (Wallet) Set Resource Management Rules

Data Pump

View Data Pump jobs and use our wizard to quickly create and run import jobs.



The screenshot displays the Data Pump interface. On the left, a sidebar shows a list of jobs, including one named 'IMP_SO_31776-14_15_43' with a green checkmark indicating success. The main area on the right shows a 'Wizard' for creating a new job. The wizard has a progress bar with steps: Create, Import, Run, Shipping, Cleanup, and Summary. The 'Create' step is currently active. Below the progress bar, there are fields for 'Job Name' (set to 'imp_so_31776-14_15_43') and 'Job Type' (set to 'Import'). A table below these fields shows the job's status, with columns for 'Job Name', 'Job Type', and 'Status'. The 'Status' column shows 'SUCCEEDED'.

Job Name	Job Type	Status
imp_so_31776-14_15_43	Import	SUCCEEDED

[Documentation](#)

Open

 Performance Hub Database Dashboard

Performance Hub

The Performance Hub page shows performance data for a specified time period, including SQL Monitoring reports and Active Session History (ASH) Analytics.

[Documentation](#)[Open](#)

 Performance Hub Database Dashboard

Database Dashboard

Monitor database activity charts such as CPU usage, number of executing SQL statements, and wait events formerly found on your *Autonomous Database Service Console*.




[Documentation](#)

Open



You can easily connect SQL Developer for VS Code to Autonomous Database



```
select cloud_identity from v$pdb;
```

```
{  
  "DATABASE_NAME" : "MYADB1",  
  "INFRASTRUCTURE" : "Serverless",  
  "SERVICE" : "ADW",  
  "APPLICATIONS" : [ "GRAPH", "ODI", "ORDS", "DATABASEACTIONS",  
                     "OMLMOD", "APEX", "OML" ],  
  "AUTOSCALABLE_STORAGE" : false,  
  "COMPUTE_AUTOSCALING" : true,  
  ...  
}
```



You can use Autonomous Database as RMAN Recovery Catalog database

- [Details](#)



ADB Features

Some of these are very important



Tablespace Auto-Remapping

Dictionary Managed



Locally Managed

Manual Space Allocation



ASSM

Uniform Extent Allocation



Auto Allocation

Smallfile / Bigfile Tablespace Conversion

Original DDL:

```
CREATE TABLESPACE tbs DATAFILE '+DATA/mydb/tbs.dbf' ...
```

Converted DDL:

```
CREATE BIGFILE TABLESPACE tbs ...
```

- **ADB-D** transparently creates a **Bigfile tablespace** with no application changes required
- **Datafile name is omitted:**
 - Oracle Managed Files (OMF)
 - File naming and management handled internally



Automatic Indexing

Indexes

Indexes are critical for SQL performance

Manual index creation:

- Requires solid understanding of schema and queries
- Quickly becomes outdated as workloads change

Poor index choices = **performance issues**

Automatic Indexing

- It's an optional feature
- You can use it in **Oracle Autonomous Database**
- **Disabled** by default
- Continuously monitors workload
- Automatically:
 - Identifies indexing opportunities
 - Creates and drops indexes as needed
 - Maintains and adapts over time

Automatic Indexing | Enabling

1. Enable it with DBMS_AUTO_INDEX.CONFIGURE procedure

```
EXEC DBMS_AUTO_INDEX.CONFIGURE( 'AUTO_INDEX_MODE' , 'IMPLEMENT' );
```

- This enables automatic indexing in a database
- It creates any new auto indexes as visible indexes, so that they can be used in SQL statements

Automatic Indexing | Enabling

2. Use the DBMS_AUTO_INDEX package to report on the automatic task and to set automatic indexing preferences
 - When automatic indexing is enabled, index compression for auto indexes is enabled by default

For JSON data there may be restrictions for Automatic Indexing in ADB

Automatic Indexing | Some options

Automatic indexing can be initiated in a *report-only* mode:

```
EXEC DBMS_AUTO_INDEX.CONFIGURE('AUTO_INDEX_MODE','report only');
```

```
BEGIN
    :rep := dbms_auto_index.report_activity(
        activity_start => sysdate-2,
        section        => 'ALL',
        level           => 'ALL',
        type            => 'HTML');
END;
/
```

Automatic Indexing | Lists

Inclusion list:

```
begin
  dbms_auto_index.configure(
    parameter_name => 'auto_index_schema',
    parameter_value => 'HR',
    allow          => TRUE);
end;
/
```

Exclusion list:

```
begin
  dbms_auto_index.configure(
    parameter_name => 'auto_index_schema',
    parameter_value => 'HR',
    allow          => FALSE);
end;
/
```

Parameters

```
select parameter_name, parameter_value from dba_auto_index_config order by 1;
```

ry Result x

  SQL | All Rows Fetched: 8 in 0.075 seconds

PARAMETER_NAME	PARAMETER_VALUE
AUTO_INDEX_COMPRESSION	ON
AUTO_INDEX_DEFAULT_TABLESPACE	(null)
AUTO_INDEX_MODE	OFF
AUTO_INDEX_REPORT_RETENTION	373
AUTO_INDEX_RETENTION_FOR_AUTO	373
AUTO_INDEX_RETENTION_FOR_MANUAL	(null)
AUTO_INDEX_SCHEMA	(null)
AUTO_INDEX_SPACE_BUDGET	100

Further Information: Auto Indexing

<https://docs.oracle.com/en/cloud/paas/autonomous-database/serverless/adbsb/autonomous-experienced-database-features.html>

Manage Automatic Indexing on Autonomous Database

<https://docs.oracle.com/en-us/iaas/autonomous-database-serverless/doc/autonomous-auto-index.html>

<https://docs.oracle.com/en/cloud/paas/autonomous-database/serverless/adbsb/autonomous-auto-index.html>

Manage Automatic Partitioning on Autonomous Database

<https://docs.oracle.com/en-us/iaas/autonomous-database-serverless/doc/mdw-managing-automatic-partitioning-autonomous-database.html>



Automatic Partitioning



Automatic Partitioning

Automatically analyzes workload

Partitions tables and indexes

Improves performance and manageability of large tables

Transparent to users — no manual action required

Complements manual partitioning; doesn't override it

Partitioning Methods Used

AUTOMATIC INTERVAL

- Best for range-based values

LIST AUTOMATIC

- For distinct partition key values

HASH

- Based on hash of partition key

Only supports single-column, single-level partitioning

How Automatic Partitioning Works

1. Analyze workload
 2. Identify candidate tables
 3. Evaluate partitioning strategies
 - Estimates IO benefits
 - Verifies performance improvements
 4. Recommend or implement optimal strategy
- Only runs **on demand** via PL/SQL API `DBMS_AUTO_PARTITION.RECOMMEND_PARTITION_METHOD`
 - This is **not** a background task such as auto-indexing

How Automatic Partitioning Works

Enable, and **implement** recommendations

```
EXEC DBMS_AUTO_PARTITION.CONFIGURE('AUTO_PARTITION_MODE', 'IMPLEMENT');
```

Enable recommendations, but **do not implement** those recommendations

```
EXEC DBMS_AUTO_PARTITION.CONFIGURE('AUTO_PARTITION_MODE', 'REPORT ONLY');
```

You can use `AUTO_PARTITION_SCHEMA` and `AUTO_PARTITION_TABLE` settings for some fine tuning

Generate and apply recommendations

Apply **automatically** for eligible tables

```
EXEC DBMS_AUTO_PARTITION.CONFIGURE('AUTO_PARTITION_MODE', 'IMPLEMENT');
```

```
VARIABLE RECOMMENDATION_ID VARCHAR2(32);  
BEGIN  
    :RECOMMENDATION_ID := DBMS_AUTO_PARTITION.RECOMMEND_PARTITION_METHOD();  
END;  
/
```

```
SELECT TABLE_OWNER, TABLE_NAME, PARTITION_METHOD, PARTITION_KEY  
FROM DBA_AUTO_PARTITION_RECOMMENDATIONS
```

```
SELECT DBMS_AUTO_PARTITION.REPORT_LAST_ACTIVITY() FROM DUAL;
```

Generate and apply recommendations

Apply **manually** for eligible tables

```
EXEC DBMS_AUTO_PARTITION.CONFIGURE('AUTO_PARTITION_MODE','REPORT ONLY');
```

```
VARIABLE RECOMMENDATION_ID VARCHAR2(32);  
BEGIN  
    :RECOMMENDATION_ID := DBMS_AUTO_PARTITION.RECOMMEND_PARTITION_METHOD();  
END;  
/
```

```
SELECT TABLE_OWNER, TABLE_NAME, PARTITION_METHOD, PARTITION_KEY  
       FROM DBA_AUTO_PARTITION_RECOMMENDATIONS
```

```
BEGIN  
    DBMS_AUTO_PARTITION.APPLY_RECOMMENDATION(RECOMMENDATION_ID => :RECOMMENDATION_ID);  
END;  
/
```

Generate and apply recommendations

Apply **manually** for a **single table**

```
EXEC DBMS_AUTO_PARTITION.CONFIGURE('AUTO_PARTITION_MODE','REPORT ONLY');
```

```
VARIABLE RECOMMENDATION_ID VARCHAR2(32);  
BEGIN  
    :RECOMMENDATION_ID := DBMS_AUTO_PARTITION.RECOMMEND_PARTITION_METHOD(  
        TABLE_OWNER => 'TPCH',  
        TABLE_NAME  => 'LINEITEM');  
END;  
/
```

```
SELECT TABLE_OWNER, TABLE_NAME, PARTITION_METHOD, PARTITION_KEY  
FROM DBA_AUTO_PARTITION_RECOMMENDATIONS
```

```
SELECT REPORT  
FROM DBA_AUTO_PARTITION_RECOMMENDATIONS  
WHERE RECOMMENDATION_ID = :RECOMMENDATION_ID;
```


Generate and apply recommendations

Apply **manually** for a **single table**

```
BEGIN
    DBMS_AUTO_PARTITION.APPLY_RECOMMENDATION(
        RECOMMENDATION_ID => :RECOMMENDATION_ID);
END;
/
```

If you want to apply the recommendation to a different database:

```
SELECT MODIFY_TABLE_DDL
FROM DBA_AUTO_PARTITION_RECOMMENDATIONS
WHERE RECOMMENDATION_ID = :RECOMMENDATION_ID;
```



Fast Ingest

Using Fast Ingest on Oracle ADB

- Optimized for **high-frequency, single-row inserts**
- Ideal for **IoT** and "**fire and forget**" workloads
- Improves **insert throughput**
- Not designed for **full ACID compliance**

Example use cases:

- Sensor data
- Smart meter readings
- Traffic cameras

When to Use Fast Ingest

Use it when:

- Ingesting **millions of small events per minute**
- ACID is **not mandatory**
- **Post-processing** is acceptable for cleanup/validation
- You want **best-in-class ingest performance**

Avoid if:

- You need **real-time durability**
- Workload is **OLTP-style transactions**

Prerequisites for Fast Ingest

1. Enable Optimizer to use hints

- Set `optimizer_ignore_hints=FALSE` at session or system level
- Required to allow use of the Fast Ingest hint

Prerequisites for Fast Ingest

2. Create an Eligible Table

- Table must not have unsupported features ([see documentation](#))
 - E.g., certain constraints, triggers, and features disqualify fast ingest usage
- Some unsupported features are lifted in 23ai

3. Hint

- Enable fast ingest for inserts by specifying the MEMOPTIMIZE_WRITE hint in the INSERT statement

Using Fast Ingest

```
CREATE TABLE test_fast_ingest (  
  id NUMBER PRIMARY KEY,  
  test_col VARCHAR2(15)  
)  
SEGMENT CREATION IMMEDIATE  
MEMOPTIMIZE FOR WRITE;
```

```
INSERT /*+ MEMOPTIMIZE_WRITE */  
  INTO test_fast_ingest  
  VALUES (1, 'test');
```

Further Information

Use Fast Ingest on Autonomous Database

<https://docs.oracle.com/en/cloud/paas/autonomous-database/serverless/adbsb/autonomous-fast-ingest.html>



Statistics



Autonomous Database Dedicated gathers stats like any other Exadata database

```
select  origin, start_time
from    dba_auto_stat_executions
order by start_time;
```

ORIGIN

START_TIME

AUTO_TASK

19-JUN-25 10.00.02.475575000 PM GMT

...

HIGH_FREQ_AUTO_TASK

20-JUN-25 08.35.38.149345000 AM GMT

HIGH_FREQ_AUTO_TASK

20-JUN-25 08.50.38.568564000 AM GMT

HIGH_FREQ_AUTO_TASK

20-JUN-25 09.05.39.171647000 AM GMT

HIGH_FREQ_AUTO_TASK

20-JUN-25 09.20.39.558415000 AM GMT

```
select  operation, target, start_time, job_name
from    DBA_OPTSTAT_OPERATIONS
order by start_time;
```

OPERATION	TARGET	START_TIME	JOB_NAME
gather_database_stats (auto)	AUTO	19-JUN-25 10.00.02.295346000 PM GMT	ORA\$AT_OS_OPT_SY_349
purge_stats		19-JUN-25 10.00.50.552518000 PM GMT	ORA\$AT_OS_OPT_SY_349
purge_stats		19-JUN-25 10.00.51.783925000 PM GMT	ORA\$AT_OS_OPT_SY_349
...			
gather_database_stats (auto)	AUTO	20-JUN-25 08.13.29.267822000 AM GMT	ORA\$_ATSK_AUTOSMT
gather_database_stats (auto)	AUTO	20-JUN-25 08.28.29.846908000 AM GMT	ORA\$_ATSK_AUTOSMT
gather_database_stats (auto)	AUTO	20-JUN-25 08.43.30.349803000 AM GMT	ORA\$_ATSK_AUTOSMT
gather_database_stats (auto)	AUTO	20-JUN-25 08.58.30.842146000 AM GMT	ORA\$_ATSK_AUTOSMT
gather_database_stats (auto)	AUTO	20-JUN-25 09.13.31.355042000 AM GMT	ORA\$_ATSK_AUTOSMT
gather_database_stats (auto)	AUTO	20-JUN-25 09.28.31.933039000 AM GMT	ORA\$_ATSK_AUTOSMT
gather_database_stats (auto)	AUTO	20-JUN-25 09.43.32.600796000 AM GMT	ORA\$_ATSK_AUTOSMT
gather_database_stats (auto)	AUTO	20-JUN-25 09.58.32.941213000 AM GMT	ORA\$_ATSK_AUTOSMT





Autonomous Database Serverless introduces new ways of gathering stats



Data Warehouse does not use
Automatic Statistics Gathering

Data Warehouse Serverless

- In DWH after loading, DMLs are atypical or infrequent
 - Gathering stats on such a workload is waste of resources
 - Automatic Statistics Gathering is disabled
- Gathers statistics automatically following most direct load operations
 - If you use conventional DML, you must gather statistics manually



Transaction Processing uses the regular Automatic Statistics Gathering job

Dictionary Views

- The DBA_OPTSTAT and DBA_AUTOTASK views are disabled in Autonomous Database Serverless
- Use other dictionary views for information on stats gathering, like
 - DBA_TAB_STATISTICS
 - DBA_TAB_COL_STATISTICS
 - DBA_IND_STATISTICS



High-frequency Statistics Gathering complements Automatic Statistics Gathering

High-frequency Statistics Gathering

- Lightweight task that only gathers stale statistics
- Reduce the period where an objects might have stale statistics
- Helps avoid cardinality misestimates on volatile tables

```
select
  dbms_stats.get_prefs ('AUTO_TASK_STATUS') high_freq,
  dbms_stats.get_prefs ('AUTO_TASK_INTERVAL') run_interval,
  dbms_stats.get_prefs ('AUTO_TASK_MAX_RUN_TIME') max_run_time
from dual;
```

HIGH_FREQ	RUN_INTERVAL	MAX_RUN_TIME
ON	900	3600

--Shows the last result of auto stats gathering

--Works in ADB-S also

```
select status, elapsed_time, last_schedule_time
from   dba_autotask_schedule_control
where  task_name='Auto Statistics Management Task';
```

STATUS	ELAPSED_TIME	LAST_SCHEDULE_TIME
SUCCEEDED	13	2025-06-20T10:05:34.942Z



Real-Time Statistics updates the most essential stats during DML

- Disabled by default

Real-Time Statistics

- Augments traditional statistics collection
- When a DML operation is currently modifying a table, the most essential statistics are computed
- Does not gather statistics, but updates existing statistics, like
 - Number of rows
 - Number of table blocks
 - High/low values

--Control the use of Real-Time Statistics using a parameter
--Available in ADB-S and ADB-D, but disabled by default

```
alter system set optimizer_real_time_statistics=true;
```



Automatic Statistics Gathering also updates stale dictionary stats

- Otherwise, file a service request



Autonomous Database Serverless creates histograms on all columns

- Autonomous Database Dedicated uses the default **FOR ALL COLUMNS SIZE AUTO**

```
--If stats gathering takes too long on a specific table due to histograms  
--you can disable histograms using a table-specific preference
```

```
exec dbms_stats.set_table_prefs(  
    ownname => 'SCOTT',  
    tabname => 'EMP',  
    pname   => 'METHOD_OPT',  
    pvalue  => 'FOR ALL COLUMNS SIZE 1');
```

Further Information

- [Manage Optimizer Statistics on Autonomous Database Serverless](#)
- [High Performance Features in Autonomous Database on Dedicated Exadata Infrastructure](#)
- [How to Use High-frequency Statistics Gathering](#)
- [Configuring High-Frequency Automatic Optimizer Statistics Collection](#)
- [Real-Time Statistics](#)



Performance



Autonomous Data Warehouse ignores optimizer hints - including **PARALLEL**

- Applies to Serverless and Dedicated

```
select name, value, isses_modifiable
from   v$parameter
where  name like 'optimizer_ignore%';
```


NAME	VALUE	ISSES_MODIFIABLE

optimizer_ignore_hints	TRUE	TRUE
optimizer_ignore_parallel_hints	TRUE	TRUE




SQL Plan Management increases plan stability


SQL Plan Management



SQL Plan Management



SQL Plan Management






Autonomous Database Dedicated uses Automatic SQL Plan Management

- With background verification

SQL Plan Management






Autonomous Database Serverless uses Real-time SQL Plan Management

- You can enable it in Autonomous Database Dedicated

SQL Plan Management





Every Autonomous Database uses Automatic SQL Tuning Sets



Effective use of SQL Plan Management
requires an application using binds



You can still use
SQL patches and SQL profiles

1 SQL Patch

overrides

2 SQL Plan Baseline

overrides

3 SQL Profile




Use Performance Hub to monitor ongoing activity

- Found in Database Actions

Quick Select Time Range Time Zone

Last Hour Jun 27, 2025, 8:19:41 AM - 9:19:41 AM Browser (GMT+02:00) Hide Activity Summary Refresh

Activity Summary (Average Active Sessions) ⓘ



ASH Analytics SQL Monitoring

Applied Filters ⓘ

Top 100 by Last Active Time SQL Duration ⓘ min to max Filter ⓘ Filter by Status, SQL ID or User Name

Terminate Session

Status	Duration	SQL ID	SQL Plan Hash	User Name	Parallel	Database Time	I/O Requests	SQL Text
●	1.80 min	6rgmagaf4vb6y	658468918	ADMIN		1.79 min	22	select * from (select q_*, ...
●	0.19 sec	dqbb3qm0f3dhm	1081506878	ADMIN	2	0.19 sec		select object_name from sys.al...
●	1.00 sec	8j7ta7nyd265m	1081506878	ADMIN	2	0.05 sec		select object_name from sys.al...
●	0.03 sec	b434n54up4hr7	1081506878	ADMIN	2	0.03 sec		select procedure_name from sys...
●	0.04 sec	d9a710xj6xk91	1081506878	ADMIN	2	0.04 sec		select object_name from sys.al...
●	0.41 sec	6xwdfvna5d81	3649520969	ADMIN	2	0.41 sec	14	SELECT max(owner ' obj...

Quick Select

Time Range

Time Zone


Last Hour

Jun 27, 2025, 8:19:41 AM - 9:19:41 AM

Browser (GMT+02:00)

Hide Activity Summary

Refresh



Applied Filters ⓘ none

Top 100 by Last Active Time

SQL Duration ⓘ min to max

Filter ⓘ Filter by Status, SQL ID or User Name

Terminate Session

Status	Duration	SQL ID	SQL Plan Hash	User Name	Parallel	Database Time	I/O Requests	SQL Text
●	1.80 min	6rgmagaf4vb6y	658468918	ADMIN		1.79 min	22	select * from (select q_*, ...
●	0.19 sec	dqbb3qm0f3dhm	1081506878	ADMIN	2	0.19 sec		select object_name from sys.al...
●	1.00 sec	8j7ta7nyd265m	1081506878	ADMIN	2	0.06 sec		select object_name from sys.al...
●	0.03 sec	b434n54up4hr7	1081506878	ADMIN	2	0.03 sec		select procedure_name from sys...
●	0.04 sec	d9a710kz6xk91	1081506878	ADMIN	2	0.04 sec		select object_name from sys.al...
●	0.41 sec	6xwdfvwa5d81	3649520969	ADMIN	2	0.41 sec	14	SELECT max(owner ' ' obj...

Quick Select

Time Range

Time Zone


Last Hour

Jun 27, 2025, 8:19:41 AM - 9:19:41 AM

Browser (GMT+02:00)

Hide Activity Summary

Refresh



ASH Analytics

SQL Monitoring

Applied Filters ⓘ

Top 100 by Last Active Time

SQL Duration ⓘ min to max

Filter ⓘ Filter by Status, SQL ID or User Name

Terminate Session

Status	Duration	SQL ID	SQL Plan Hash	User Name	Parallel	Database Time	I/O Requests	SQL Text
	1.80 min	6rgmagaf4vb6y	658468918	ADMIN		1.79 min	22	select * from (select q_*, ...
	0.19 sec	dqbb3qm0f3dhm	1081506878	ADMIN	2	0.19 sec		select object_name from sys.al...
	1.00 sec	8j7ta7nyd265m	1081506878	ADMIN	2	0.05 sec		select object_name from sys.al...
	0.03 sec	b434n54up4hr7	1081506878	ADMIN	2	0.03 sec		select procedure_name from sys...
	0.04 sec	dha710kzfbk91	1081506878	ADMIN	2	0.04 sec		select object_name from sys.al...
	0.41 sec	6xadfvvta5d81	3649520969	ADMIN	2	0.41 sec	14	SELECT max(owner ' ' obj...

← Back Real-time SQL Monitoring for SQL 6rgmagaf4vb6y

Time Zone: Browser (GMT+02:00)

Save Report

Refresh

Overview

General

Status: Executing

SQL Text: select * from (select q_*, row_number() over ...

Execution Started: Jun 27, 2025, 9:17:53 AM GMT+02:00

Last Refresh Time: Jun 27, 2025, 9:21:45 AM GMT+02:00

Execution ID: 83886080

User Name: ADMIN@LJITYZRCWFWPE_SECONDADB

Fetch Calls: 0

Time & Wait

Duration: 3.87 m

Database Time: 3.86 m

PL/SQL & Java: 0 s

Wait Activity: 0%

I/O

Buffer Gets: 332

I/O Requests: 22

I/O Bytes: 2.4 MB

Plan Statistics | SQL Text | Activity | Metrics

Plan Hash Value: 658468918 | Plan Size: 14

View Option: Tabular Execution Plan | Plan Notes

Operation	Object	Info	Line	Timeline	Executions	Est. Rows	Rows	Mem	Temp	I/O Requests
SELECT STATEMENT			0		1					
RESULT CACHE (SEL\$1)	4stjma7ascchfhp0i		1		1	1				
VIEW (SEL\$2)			2		1	1				
WINDOW NOSORT STOPKEY (SEL\$2)			3		1	1				
VIEW (SEL\$3)			4		1	1				
FILTER (SEL\$3)			5		1					

Plan Statistics

SQL Text

Activity

Metrics

Plan Hash Value: 658468918 Plan Size: 14

View Option ⓘ

Tabular Execution Plan

Plan Notes

Operation	Object	Info	Line	Timeline	Executions	Est. Rows	Rows	Mem	Temp	I/O Requests
✚ SELECT STATEMENT			0		1					
✚ RESULT CACHE (SEL\$1)	4trjma7aeothfhp0x		1		1	1				
✚ VIEW (SEL\$2)		19	2		1	1				
✚ WINDOW NOSORT STOPKEY (SEL\$2)		19	3		1	1				
✚ VIEW (SEL\$3)			4		1	1				
✚ FILTER (SEL\$3)		19	5		1					
✚ SORT AGGREGATE			6		1	1				
✚ MERGE JOIN CARTESIAN			7		1	377 G	0.348 M			
✚ MERGE JOIN CARTESIAN			8		1	411 K	0.305			
TABLE ACCESS STORAGE FULL (SEL\$3)	CHANNELS		9		1	5	2			
✚ BUFFER SORT			10		1	82 K	9,382	1.9 MB		
TABLE ACCESS STORAGE FULL (SEL\$3)	COSTS		11		1	82 K	82 K			5
✚ BUFFER SORT			12		0.382	919 K	4,248 M	21.1 MB		
TABLE ACCESS STORAGE FULL (SEL\$3)	SALES		13		1	919 K	919 K			1



Using SQL Trace

SQL Trace



```
EXEC DBMS_CLOUD.CREATE_CREDENTIAL( ... );  
ALTER DATABASE PROPERTY SET DEFAULT_CREDENTIAL = '...';
```



```
ALTER DATABASE PROPERTY  
SET DEFAULT_LOGGING_BUCKET='https://objectstorage...';
```



```
ALTER SESSION SET SQL_TRACE=TRUE;
```

```
CREATE OR REPLACE TRIGGER TRACE_LOGON
AFTER LOGON ON DATABASE
BEGIN
    IF SYS_CONTEXT('USERENV','SESSION_USER') = 'SH' THEN
        EXECUTE IMMEDIATE 'ALTER SESSION SET SQL_TRACE=TRUE';
    END IF;
END;
/
```

Source: <https://blogs.oracle.com/datawarehousing/post/enable-sql-trace-on-autonomous-database>

Further Reading

Hints:

Serverless: <https://docs.oracle.com/en-us/iaas/autonomous-database-serverless/doc/manage-optimizer-stats.html>

Dedicated: <https://docs.oracle.com/en/cloud/paas/autonomous-database/dedicated/hpfad/index.html#HPFAD-GUID-6C0FC667-B8E7-488C-92DF-735C1CD95536>

SQL Plan Management:

[What is automatic SQL plan management and why should you care?](#)

[What is Real-time SQL plan Management?](#)

[What is the Automatic SQL Tuning Set?](#)

[What's the Difference Between SPM Auto Capture and Auto SPM?](#)

[Automatic SQL Tuning Sets \(ASTS\) 19c RU 19.7 Onwards \(Doc ID 2686869.1\)](#)

Tracing:

Serverless: [Enable SQL Tracing on Autonomous Database](#)

Dedicated: [Use SQL Tracing on Autonomous Database on Dedicated Exadata Infrastructure](#)

[Enable SQL Trace on Autonomous Database](#)



Try it out, please!!

- We are looking for reference customers
- Get in touch with us when you tested it

Find Slides and Much More on Our Blogs



MikeDietrichDE.com

Mike.Dietrich@oracle.com



dohdatabase.com

Daniel.Overby.Hansen@oracle.com



DBArj.com.br

Rodrigo.R.Jorge@oracle.com




AlexZaballa.com

Alex.Zaballa@oracle.com



YouTube | @UpgradeNow



[Link](#)

- 300+ videos
- New videos every week
- No marketing
- No buzzwords
- All tech



Thank You

